

A HUMAN PERFORMANCE MODELLING APPROACH TO
INTELLIGENT DECISION SUPPORT SYSTEMS

MICHAEL S. MCCOY
DESIGN ENGINEERING - HUMAN FACTORS
MCDONNELL AIRCRAFT COMPANY
ST. LOUIS, MO 63166

RANDY M. BOYS
ARTIFICIAL INTELLIGENCE LABORATORY
DEFENSE SYSTEMS AND ELECTRONICS GROUP
TEXAS INSTRUMENTS, INCORPORATED
DALLAS, TEXAS 75266

ABSTRACT

Manned space operations require that the many automated subsystems of a space platform be controllable by a limited number of personnel. To minimize the interaction required of these operators, artificial intelligence techniques may be applied to embed a human performance model within the automated, or semi-automated, systems, thereby allowing the derivation of operator intent. A similar application has previously been proposed in the domain of fighter piloting, where the demand for pilot intent derivation is primarily a function of limited time and high workload rather than limited operators. The derivation and propagation of pilot intent is presented as it might be applied to the Darpa/AFWAL "Pilots Associate" or AAMRL "Super Cockpit" programs.

INTRODUCTION

Traditionally, a Human Performance Model is used as a design tool to predict various man-machine system designs. Such a model could be incorporated into an intelligent computer system for two additional applications. First, the model could be one element of a decision and task support system. It could predict the operator's forthcoming tasks, anticipate upcoming decisions, and formulate any necessary decision or execution aids. Second, the model could serve as a performance monitoring tool by analyzing the differences (and causes of differences) between expected and actual operator actions. These differences may indicate that: (1) modeled operator activities and goals were incorrect and require updating; (2) steps in a procedure were omitted by the operator; or (3) that critical information has not been presented to the operator. This data will be factored into the decision/execution support system, thus enabling the computer system to recommend or perform corrective actions.

The Pilot's Associate (PA) system is a decision and task support system which will use artificial intelligence (AI) technology to aid post-1995 fighter aircraft pilots. One of the expert systems within PA is the Pilot/Vehicle Interface (PVI) system. This system, employing machine intelligence and enhanced control and

display devices, is being designed to improve the pilot's situation awareness. The primary goal of the PVI expert is to manage information flow between the pilot and the other PA expert systems (Systems Status, Mission Planning, Situation Assessment, Tactical Planning, and the System Executive) (McCoy & Boys, 1987).

Another example of the developing need for an "Electronic Crewmember" is that proposed for the Super Cockpit (SC) program. The SC system will immerse fighter pilots in a visual, aural, and tactile world that is a combination of real-world and computer-generated events. Compared with the PA program, SC is more of a technology pull for these "virtual world" subsystems than expert system technology itself. The PA system would be a subsystem available to the pilot through the SC interface.

Critical to the performance of both programs is the ability to derive and reason from Pilot Intent (PI). This capacity is provided by the Pilot Intent subsystem of PA and the Pilot Intent Inference Engine of Super Cockpit (McCoy & Boys, 1987 and Martin 1986). PI helps shift the pilot's role from operator to system manager so that the pilot can specify that something should be done, without precisely specifying how it should be done. Thus, PI must be able to identify the activities in which the pilot is involved, their relative priorities and execution constraints, and the intent of any non-deterministic pilot commands. This knowledge is then merged with any new tasks identified by the various Electronic Crewmember's (EC) subsystems, resulting in EC's tailoring of displays and man-machine interactions.

To identify pilot intent, the PI subsystem must assume that the pilot exhibits purposive behavior. By observing the behavior of the pilot, the system can derive the purpose or intent of the pilot while performing a particular task. The pilot's intent is derived from three sources: (1) discrete commands; (2) complex commands (requiring EC decomposition); and (3) inferred intent (derived from observing pilot behavior as a function of the current context). It will be shown that this derivation of purpose from behavior is a reverse application of the Human Performance Model (HPM) method of predicting behavior based on a known purpose.

*Portions of this document have been previously published (McCoy & Boys, 1987)

Traditionally, HPM's are used to predict consequences of specific situation-response mechanisms. These models typically predict operator behavior, performance, and workload within a given situation. By modeling the operator's performance of known tasks, the consequences of assigning new tasks and increasing operator duties may be predicted. These models can also predict the effect of relieving the pilot of specific tasks through automation. The projective nature of EC mandates these applications of HPM's, but this discussion extends their utility to the real-time EC system itself.

The Electronic Crewmember EC system must provide information in a timely manner and maintain the pilot's situation awareness. Since the situation is dynamic the system must be capable of knowing the goals of the pilot and of reasoning about the decisions it must make, as well as actions it must perform, to meet those goals. One method of providing the system with this ability involves incorporating a model of the human within the EC system. This model would allow for predictions of the pilot's decisions, actions, and information needs, thus providing the system with a means of not only anticipating pilot requirements, but also using these predictions to supplement the pilot and improve overall system performance.

The efforts put into developing the traditional HPM can, therefore, apply to the embedded HPM as well, particularly if the context for this PI performance model is supplemented by the other Electronic Crewmember experts. The correspondence of these two models is presented in the pilot intent interpretation section. Once the relationship between the HPM and PI has been established, the analysis procedure, applicable to the development of both, will be presented. This analysis will be shown to provide all of the information necessary to develop the purpose-behavior relationship required for both models. Finally, applications of the combined human performance/pilot intent model are described.

HUMAN PERFORMANCE MODELS

Although the title "Human Performance Model" implies a concentration on the human's role in a man-machine system, the only way to truly model the human's performance in a system is to include a model of the machine and of the environment that imposes demands on the human. Therefore, a closed-loop system model is developed, with statistics being collected (or predicted) on the human's participation in the total system performance. Several purposes for HPM's are: (1) explanation of the system being studied; (2) analysis of the system being studied; (3) assessment of the design of a new system; and (4) prediction of performance, or operator workload, within an existing system.

As an explanatory device, HPM's can be used to: define a system or problems within the system; isolate ambiguous relationships between inputs and outputs to the system; and enhance the analyst's understanding of the underlying dynamics of the system. The resulting networks often capture the

functional intent of the system in a dynamic context, and can serve as input to training or development aids. These 'explanations' almost invariably exist at more than one level of abstraction, hence they are applicable to a number of different purposes (users, trainers, funding sources, analysts, etc.).

A second purpose for HPM's is as an analysis vehicle. When developing the model, the analyst is required to investigate and determine critical elements, components, and issues within the system being studied. In addition, the model can be used to investigate hypothetical relationships between various components of the system and new components to be added to the system. By using this method, much insight can be obtained about system characteristics without physically interfering with an operational system.

The HPM can also be used during design to assess, and aid in planning, new systems. In this way, required or suggested changes to the system design can be identified early in the design process, thus minimizing the risk of delays to the project during development. The advantages to life-cycle costs provided by 'up-front' work have been well documented.

A final application of HPM's is to predict the effects of proposed solutions to the problems being studied on the existing system (be it physically or hypothetically existent). By creating a model reflecting the proposed solution and then comparing the results to the simulation of the system containing the problem, an analysis can be made of the relative improvement on performance and workload based on this solution.

There are several critical components of the HPM which interact to reflect total system performance. They consist of: (1) system demands; (2) cognitive situation assessment; (3) decision making or task selection; and (4) task/procedure execution. Each of these components of the model will be explained in detail.

System Demands

Demands on the system are generated from several sources: (1) system dynamics; system malfunctions; (2) environmental factors; (3) situation contingencies; and (4) mission status. Each of these sources must be modeled separately. The first, system dynamics, represents a model of the machine, in this case the vehicle. This model represents the dynamics of the system, including the spatial position of the aircraft, the state of its various weapons, etc. When the aircraft takes off, one demand on the system is to retract the landing gear. Another demand is to raise flaps. These are examples of demands imposed on the operator by the system dynamics model.

Another model to be developed for the system is the system malfunction model. This model simulates any malfunctions, such as oil pressure problems, that can occur during flight and are important enough to be included in the system. In this way, the potential malfunctions to be considered by the operator can be represented and accounted for in generating demands on the operator.

The environment imposes specific demands that must be accounted for in the system. These demands can be weather related, such as wind gusts influencing the aircraft's state, or threat related, from either air or surface threats. In any event, they must be modeled as potential demands imposed on the system. Finally, the mission status can impose demands.

Any combination of the above demands may impose delays in the mission, requiring extra activities to be performed by the pilot to compensate for these delays. These activities can include minor adjustments to airspeed or major revisions of the mission route. All of these demands must be maintained as inputs to the pilot's cognitive situation assessment.

Cognitive Situation Assessment

When a new demand is generated, or an existing demand is eliminated, the situation must be reassessed. This assessment consists of two major tasks: 1) reprioritizing existing demands on the system and 2) developing a plan to eliminate one or more of the demands imposed on the system. Therefore a queue of demands exists (short-term memory) which must be ranked in order of importance. The prioritization can be represented as an algorithm in which a relative significance is assigned to each demand based on the situation (mission segment, current threats, relative altitude, etc.). Another method of prioritizing demands is through a set of production rules. This method would allow for determining the demand priority from the situation (as above), but would also allow for inferring the situation from the demand priority. Because of this dual applicability, the production rule approach may be more suitable for the PI portion of this model.

When the demands have been prioritized into their relative importance, an elimination of demands, or focusing on high priority demands for consideration, must be performed. This could be modeled using a heuristic process of elimination. When a seemingly manageable subset of all demands has been developed, plans are then developed for satisfying demands. This stage of plan development examines the high priority demands and selects candidate tasks/procedures which will contribute to satisfying one or more demands. It is likely that the same task/procedure could be selected to meet several pending demands. Likewise, one demand may likely be satisfied by several tasks/procedures. In any event, all candidate task/procedures must be examined. This process is largely one of pattern matching (demands are matched to a plan library) and does not necessarily require production rules in either the HPM or PI subfunction functions which perform this task.

Decision Making (Activity Selection)

As mentioned, any particular demand could be met by alternative tasks/procedures. Alternatively, a specific task/procedure could meet many of the demands. Therefore, an assessment of the candidate tasks/procedures must be made. A list of candidate task/procedures that lend themselves to one or more of the demands must be generated.

This list will provide the basis for selecting the next activity to be performed by the operator.

A ranking of tasks/procedures must then be performed. This ranking must take into consideration the number of demands that will be met by each task/procedure, the time required to perform the task/procedure, the resources required for performance of the task/procedure, and any other metrics of task desirability. This ranking can be performed by development of an expected net gain (Baron et. al., 1980) or production rules could be applied. Unlike the assessment stage, all alternatives will, most likely, not be investigated.

The final step in the decision making routine is to select the task/procedure that "best" meets demands or improves the situation. This selection criterion could be algorithmic, such as maximizing the expected net gain, or it could be heuristic. If it is heuristic, then a set of production rules must be determined. In addition, production rules may include task ranking. After the task/procedure has been selected, the effect on the situation, including the passage of time, must be simulated.

Task/Procedure Execution

HPM's traditionally incorporate a simulation of system activities which accounts for resource utilization as well as for the advance of time. Often, this simulation can be represented in the form of a network of activities, each activity seizing resources and consuming time. In this way, the potential bottlenecks and time delays, due to resource limitations (which includes the operator's attention, perception, cognitive and motor availability) can be incorporated into the model. In addition to accounting for time delays and resource utilization, the activities will satisfy the designated demand. When the task/procedure has been completed, the cognitive situation assessment portion of the model must be invoked.

Summary

HPM's can be created using many different conventions, but if the model is to reflect the operator's cognitive processing, the style chosen must support these generic stages. The model must also include any variability (usually stated as the appropriate statistical distribution, applied Monte-Carlo) associated with not only the situational demands, but with the human operator himself. Once a model is completed, it can be used for determining inherent faults in the system as well as optimizing system execution along any of the simulated parameters. These parameters include proposed system enhancements (improvements or new capabilities) or factors associated with operator performance (load leveling, response time, precision, etc.).

PILOT INTENT INTERPRETATION

The recognition of pilot plans and goals has not been a problem in the development of contemporary fighter aircraft. The role of the man-machine interface was simple and direct: receive and respond to pilot commands (switch or stick

inputs) and occasionally provide simple alerting, landing, targeting, or other cues. These cues and other rudimentary automated systems comprise the entirety of contemporary avionics autonomy.

The notion of an "Electronic Crewmember", a computer system which interacts with the pilot in a highly dynamic fashion through the incorporation of artificial intelligence and advanced input/output technologies, will change the nature of the man-machine interface. This change has been mandated by an explosion of information and activity in the cockpit as well as by the changing nature of ownship and threat capabilities. There is more to do and less time in which to do it. There is also little time for the pilot to train a computer system into an acceptable cognizance of his needs to the point where it can begin to assist him in his tasks. The EC must maintain the information flow bandwidth required for this task, which must be looked upon as reciprocal.

The reciprocal nature of communication acts has long been recognized, but only recently has the advance of computer system technology allowed for intelligence on both sides of the interface. The dynamic and independent nature of both the human and computer conversants requires that each party maintain a model of the other's behavior, as well as a model of the other's model of ones self! This is an area where psychology and sociology have done much research, but only recently have the designers of computer systems attempted to apply these theories to a man-machine interface (Baron, et. al., 1980; Wellens & McNeese, 1987; Martin, 1986).

An analogy to syntax and semantics as applied to linguistics can be made: only certain actions can follow other actions (syntax) as opposed to knowing what actions make sense right now (semantics). An example, as applied to the piloting domain, is knowledge that the pilot is requesting a EC subfunction because he has selected the 'Systems' menu on the display (syntax) as opposed to knowledge that he is requesting the detailed version of the engine status display because there have been indications of engine problems of which the pilot would have been aware (semantic). The identification of intent, regardless of the detail at which it is determined, will allow the PVI and other system components not only to model the pilot's upcoming activity in more detail, but also to provide pilot assistance in increasingly appropriate ways.

Of concern here are those EC functions which involve: the derivation of pilot goals; the identification of active pilot tasks and procedures; and the transmission of these goals and activities to the entirety of the EC system. These tasks are done both explicitly, using pilot commands (or 'expanded' commands), or implicitly, inferring pilot concerns and plans from the pilot's activity or inactivity in light of the current or anticipated situation as reported by the various EC subfunctions. The method for this derivation of 'unstated' intent is related to the already forwarded portrayal of HPM's.

The HPM methodology portrayed a model of the mechanism that a human employs to: (1) examine a situation; (2) infer the demands generated from the situation; (3) develop plans for meeting demands and improving the situation; (4) selecting a desired plan; and (5) executing the plan, which then feeds back to effect the situation. As presented in Figure 1, PI subfunction essentially performs the reverse process. Once a specific behavior is observed, this model must infer the demands and situation that must have caused this specific behavior. Therefore, an examination of this reverse process must be made.

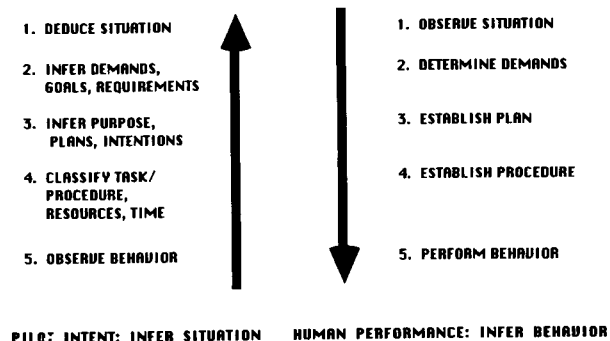


Figure 1. Operational Relationship of Pilot Intent and Human Performance Models

Behavior Observation

As the pilot operates the aircraft, he exhibits behavior that can be observed through the Hands On Throttle and Stick (HOTAS) system, for control of airframe and specific switches; the touch panel, for display change input, target designation, etc.; miscellaneous switches, for control of subsystems; Helmet Mounted Sight (HMS), for target designation, etc.; and the Voice Interactive System (VIS), for verbal inputs to the system. In addition to these command activities, other behavioral and physiological indices will be available to EC from the assumed pilot state monitoring system (used to monitor G-LOC, workload, etc. derivation). These indices include eye movements, direction of gaze, etc. When behavior has been observed it can be assumed to be purposive unless erroneous. While there is a one-to-many mapping problem at this level of analysis, the following stages successively refine the estimate of the purpose of any observed behavior.

Task/Procedure Definition

Once the input has been detected, several questions must be answered. First, does the observed behavior contribute to an existing (hypothesized) procedure with some reasonable confidence? If so, this helps confirm the belief that that procedure is being executed. This belief, in itself, may well be expressed in terms

of a probability. If the observed behavior does not confirm the belief, then there are two potential explanations. First, the believed procedure (intent) is false or second, a new procedure has been initiated in addition to those currently active.

If the task being initiated does not correspond to current expectations, yet does not conflict with them, then the task may be either an isolated task or the initiation of a new procedure to be performed in parallel with the current procedure. Knowledge must be gathered to distinguish which situation is occurring.

Purpose Inference

Given that a new task or procedure is inferred the purpose of the task or procedure must be determined. As with the HPM, a number of purposes could be met with many tasks, therefore a mechanism must be devised to distinguish when a task/procedure applies to one or another purposive behavior. The type of distinguishing characteristics necessary to determine the intention of the pilot would be current mission segment, other tasks/procedures being performed or recently completed, current threats, etc. All of this information must be maintained by the system to help infer the purpose of the behavior and may lend itself to a set of production rules. PI subsystems would be assisted by the expertise of the various other Electronic Crewmember systems in making this inference.

Demand Inference

Once intent or purpose has been inferred, the demands or goals being addressed (mapped) to the purpose can be determined. The identification of these goals is necessary if the system is to anticipate the information requirements of the pilot during execution of the procedure. In addition, given the purpose or goal for execution of a procedure, steps of the procedure are identified. While a step may be inadvertently skipped by the operator, if the step is detected by the pilot intent model, these steps can be performed for the pilot or brought to the pilot's attention depending on the level of automation. As a final case, EC may have been authorized to respond autonomously to the identified demand.

If the demands and goals have been determined, they can be used to infer the situation confronting the pilot. This inference can be used as a check in which the existing goals and demands can be compared with inferred ones, verifying current belief of the pilot's intentions. If this belief is confirmed, then the system proceeds to operate as a decision support system. If this belief is not supported, an investigation must establish a proper belief in pilot intent.

PILOT/VEHICLE INTERFACE ANALYSIS AND DESIGN APPROACH

The general methodology being followed by the MCAIR/TI team in the development of the PVI is presented in Figure 2 (McCoy, Boys, 1987). The "System Analysis" effort has only recently been completed, with current efforts being directed to various "System Design" activities. The PVI

analysis process consists of: Function Decomposition and Task Analysis. The PVI design process consists of: Task Analysis; Information Requirements Analysis; Static Display Definition; and the Pilot Intent/Human Performance Model sections (i.e., Human Performance Analysis, Dynamic Time Line Analysis, Pilot Intent Model Development, Timing Requirement Analysis, and Development of Automation Requirements). The final step in the process is to use the HPM in evaluating such procedures as: dynamic allocation of tasks; levels of autonomy (pilot/system operational relationship); load leveling of information; and those more traditional applications discussed earlier. The following paragraphs contain descriptions of each activity. It is important to remember that this analysis is iterative. Subsequent progress will be documented as necessary.

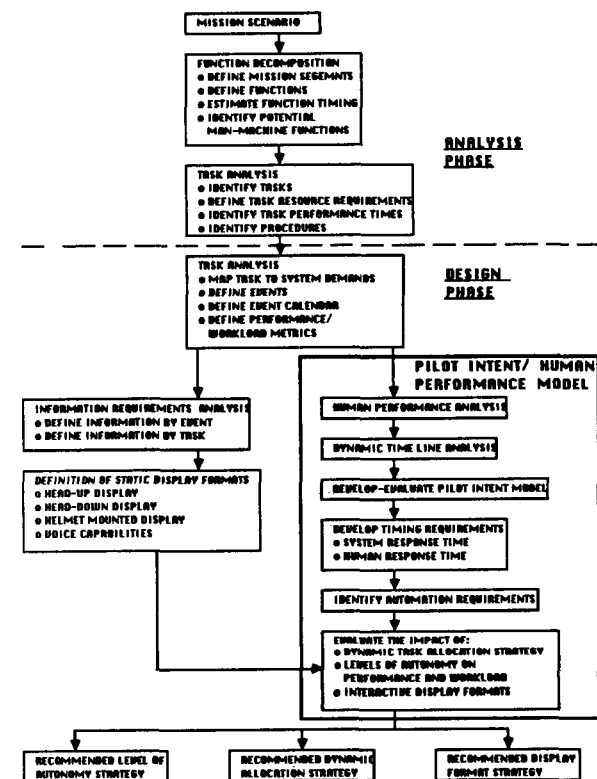


Figure 2. PVI Analysis and Design Approach

Analysis Process

Function Decomposition - The top-level mission functions in which the EC is involved (which include no less than all piloting activities) have been identified. These functions have been decomposed into subfunctions necessary to meet mission goals of survivability and effectiveness, estimated in terms of function timing requirements.

Tasks are initially classified as belonging to one of five categories: 1) tasks that can be performed only by the pilot, 2) tasks that can be performed by the pilot or a computer (the EC or conventional) but which the pilot does better, 3) tasks that can be performed by either the pilot or the computer, 4) tasks which either can perform but which the computer does better, and 5) tasks which only the computer can perform. It is sometimes necessary to make these assignments at the task level.

With the continued increase in computer capabilities (especially in the AI technology area), these latter categories continue to grow. Therefore, assumptions must be made about the level of potential automation available in the future aircraft. Completion of function decomposition leads directly to task analysis.

Task Analysis - In task analysis, individual tasks are defined to perform the functions. These tasks are performed throughout the mission and can be separated into three groups: perceptual, cognitive, and motor. The current task list is based on an analysis of a traditional high performance fighter aircraft (F-15, F/A-18) and does not reflect the task analysis or task allocation projected for an advanced combat aircraft equipped with the EC system. Many of the motor portions of tasks could be performed by various entities of the EC system. This leads to the next task required in task analysis, definition of resource requirements.

Once tasks have been defined, the resources necessary to perform the task must be determined. The term "resources" refers to a human, a machine, or a combination of the two. The definition of resource requirements is a logical extension of the function classification process. As resources are assigned to a task, the timing of that task can be determined. The timing of the task may be different given different resources. Analyses can be performed to develop the tradeoffs associated with automating tasks, varying degrees of pilot/system interaction, task allocations, etc.

Design Process

Task Analysis - The next set of activities occur simultaneously. First, a definition of events in representative missions must be made. An example of an event is detection of a threat. Given this event, what tasks are necessary to survive the threat? By answering this question, a link is established between demand for tasks and the tasks necessary to meet that demand. In addition, alternative tasks can be determined, leading to an analysis of what conditions lead to which tasks.

As tasks are defined, situations will arise in which specific tasks occur sequentially. This sequential occurrence of tasks is called a procedure. Procedures must be defined to include the conditions under which they will occur. Defining procedures allows for the summarization of tasks under specific conditions. These specific conditions will aid in determining the procedure being performed when the behavior is observed. Of course, variability of procedures (particularly when the EC system attempts to 'internalize' these

systems for dynamic use) must also be represented. The variability of procedures is represented by alternative task orderings within a procedure.

The final major activity to be addressed in the task analysis is to determine the performance and workload metrics to be used in analyzing various operating practices. There are two aspects to performance and workload metrics. First, metrics are needed to predict performance under different operating conditions. Second, metrics are used to measure performance and workload during experiments and during flight. These metrics should be developed for each task or procedure. Once this activity is complete, the task analysis is complete for both analysis and design and parallel activities related to the Information Requirements Analysis and the development of the Pilot Intent/Human Performance Models can begin.

Information Requirements Analysis - When the task analysis has progressed sufficiently, an information requirements analysis must be initiated. Preliminary versions of this analysis were conducted for the analysis stage task analysis, but they must now be mapped to the emerging EC task analysis. In addition, this information analysis should be performed for each event. This information requirements analysis will allow the EC to anticipate the information needs of the pilot throughout the mission, thereby contributing to the knowledge engineering effort. A major consideration is the timing requirements for this information. The timing requirements must be discussed in conjunction with the Pilot Intent/Human Performance Analysis.

Human Performance Analysis - The Human Performance analysis produces the HPM. The HPM must have a representation of the response characteristics of the airframe and subsystems, including the Electronic Crewmember. This representation will dictate demand for tasks to be performed during the course of the simulated scenarios. The aircraft model need only be detailed enough to account for behavior characteristics and response times given a specific control input.

The second portion of the HPM is a library of procedures and tasks. This library must be developed using the task analysis, and can be viewed as knowledge of the pilot's activities or potential activities during the mission. Alternative representations of tasks may be developed if a task can be performed by different resources, human and computer.

The final element of the HPM is the control system. This control system will scan the demands imposed by the aircraft and subsystem model, determine the highest priority demand and select the appropriate tasks or procedures to be executed. When the task or procedure is executed, it will eliminate or satisfy the appropriate task demand and allow the control structure to choose the next task, given the next highest priority demand. Thus the control model helps close the loop, representing a truly reactive system. The output of this model will reflect levels of the metrics developed in the task analysis. Examples

of the output would be resource utilization and delay time due to resource scarcity. These metrics will be used for all subsequent analysis.

Pilot Intent Model Development/Evaluation - As described earlier, the pilot intent model can be viewed as a reverse application of the HPM. The HPM must be developed to predict human activity under various circumstances throughout the mission. When this model is validated, it will be capable of predicting information demand through its identification of required tasks and procedures. By experimenting with various pilot intent strategies, this model will allow an evaluation of the appropriate strategy to be employed in the Pilot/Vehicle Interface. This design stage would be optional if it were not for the stated overlap between the HPM and PI. It is not necessarily a requirement for the next activity to be initiated.

Response Time Requirements Development - Once the human performance model has been developed, it can be used to investigate the necessary response time of the system. Given an event, such as detection of a threat, and the reaction times for the pilot, the response time of the system can be modeled. Knowing what the average allowable time for detection of a threat to its imminent impact, Monte-carlo studies can be performed to determine whether the system response time is adequate under all conditions.

Identify Automation Requirements - As a consequence of the system performance parameters gathered from the modeling efforts, suggested man-machine interactions will be identified as optimal. While this can be seen as identifying automation requirements, it actually only places lower bounds on the automation requirements. At this point, it is possible to evaluate the stabilized man-machine system in terms of acceptable variations in system configuration, man-machine interface, and performance. Referring to Figure 1, this includes an analysis of dynamic task allocation strategies, interactive modifications to the fixed display formats, and the tone of the man-machine operational relationship (level of autonomy, command modes, etc.).

APPLICATIONS OF PILOT INTENT/HUMAN PERFORMANCE MODEL

Applications Outside an Operational PVI

As was described in the discussion of HPM's, there are four major applications for the HPM. These applications include: developing response time requirements for the system to formulate information and present it to the pilot, investigating operational relationship strategies between the pilot and EC, investigating Dynamic Allocation strategies within each operational relationship, and developing strategies for load leveling the information flow to the pilot in order to prevent perceptual and cognitive saturation. By predicting the performance and workload of the pilot under various man-machine interaction strategies, an acceptable design can be chosen.

Applications Within the Operational PVI

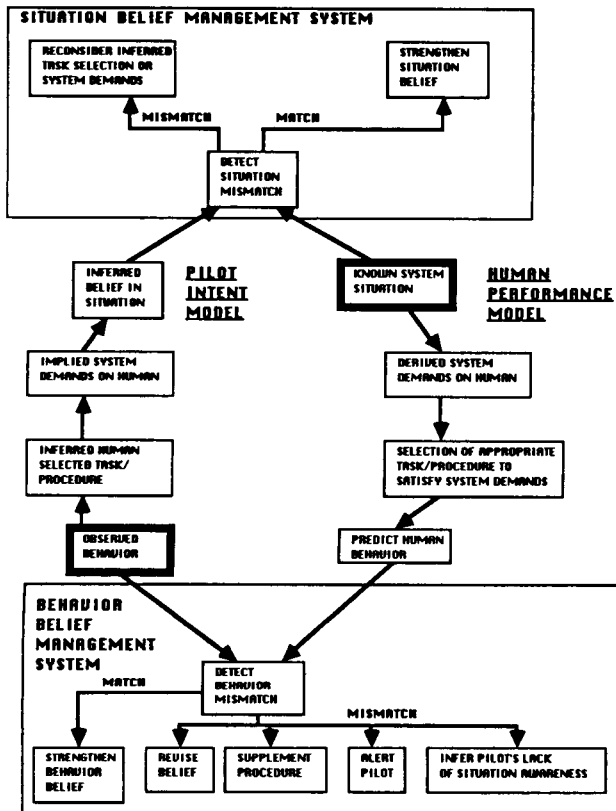
Throughout a mission, two categories of pilot intent dictate the pilot's activities. The first includes goals that are applied continuously. These goals include mission specific purpose derived during mission planning, such as the intention to deliver ordnance. In addition, this category includes more general goals which apply to all missions, such as the intention to maintain a safe flight vector, to optimize the use of expendables, etc. The second category includes situation specific goals.

There are two ways for the EC system to detect situation specific goals. First, there are overt commands to the system, such as HOTAS input or high-level commands such as 'fence check', which cues the EC to perform a sequence of activities. The second method, a more covert means of detecting intent, is performed by inferring intent from observed behavior. We have presented a method for inferring intent, and the current situation, from observed behavior. Once these intentions, overt and covert, are made known to the EC system, appropriate decision and task aids can be formulated.

The 'associate' paradigm is the motivation for the EC's requirement to be cognizant of pilot intentions. An associate can notify, advise, assist, or execute activities with, or even for, the pilot. The degree of activity or autonomy with which the EC performs these various tasks will be the product of the system's capabilities as well as the pilot's momentary desires, expectations, and workload. A robust, contextual intent interpreter can direct the the EC system in executing the associate paradigm.

An integral component of the associate paradigm is adaptive aiding, which is the system's response to changes in the mission environment, pilot workload, and pilot intentions. The dynamic allocation of tasks to the pilot or EC, along with other aspects of adaptive aiding as implemented through the man-machine interface (shared activities, sequential behavior, etc.), will ensure that the EC system is always in step with the pilot's expectations and activities. This involves much more than the storage and identification of procedures. Operator-machine interactions will be a unique product of the current situation and the operator's current desires. At its most extreme, EC's knowledge of the pilot's goals allows the EC to anticipate the activity required of it, thereby making the EC operate in a more proactive, rather than reactive, manner.

The embedded HPM within PI subsystem will provide much of the knowledge necessary to maintaining the EC's awareness of the pilot's information requirements and task assistance expectations. In addition, it provides a method for applying this knowledge toward responding to the pilot and satisfying his goals. Figure 3 describes the use of the PI/HPM model to coordinate known pilot goals, pilot behavior, and mission/environment changes reported by the various Electronic Crewmember components.



The application of the model is triggered by two external events. The first event is observed operator behavior. If the observed behavior matches the expected behavior, then the behavior belief is strengthened (and therefore the system's model of the pilot). If the expected behavior does not match the observed behavior, then there are several possible responses to include: revision of behavior belief; supplementation of pilot's errors of omission in a procedure; alerting the pilot of an error of commission (procedure definitely does not meet known situation); or infer a lack of pilot situation awareness. Observed behavior is also used by PI subsystem to infer any changes in the current situation, which, as shown below, also triggers subsequent processing.

Changes in the known situation are the second trigger for system activity. As changes in the situation are reported by the various EC expert systems, matching is again done, this time against the PI subsystem generated estimate of the situation derived from activity. A match between the estimate and reality strengthens the PI subsystem's belief in the estimated situation. A mismatch indicates a need to reconsider any inferred pilot task selection or system demands. Situation changes reported by the EC subsystems are also propagated through the PI/HPM subsystem to update expected pilot behavior.

While 'observed behavior', and 'known system situation' are the only external inputs that cause PI subsystem activity, the changes that these inputs can cause, through the reasoning capabilities of the PI/HPM subsystems, to 'inferred belief in situation' or 'predict human behavior' can also trigger the matching process.

CONCLUSION

The Electronic Crewmember system will apply new technologies to the man-machine interface of future fighter aircraft. Within the man-machine interface, a requirement has been demonstrated for the system to be cognizant of the pilot's behavior and intentions. The Pilot Intent Sub-system satisfies this requirement by applying a combined pilot intent/human performance model. This machine cognizance of the pilot provides an enhanced decision and task support capability to the EC system. This creates an 'associate' paradigm approximating that of a human associate.

REFERENCES

- Baron, S., Zacharias, G., Muralidharan, R., and Lancraft, R., "PROCURU: A model for analyzing flight crew procedures in approach to landing," Proceedings of Eight IFAC Work Congress, Tokyo, Japan, Vol. XV, pp. 71-76. (1980). (Also, NASA Report No. CR152397).
- Martin, W. L., An Assessment of Artificial Intelligence and Expert Systems Technology for Application to the Management of Cockpit Systems, AAMRL-TR-87-040. (1986)
- McCoy, M. S., Boys, R. M., "Human Performance Models Applied to Intelligent Decision Support Systems," Proceeding of NAECON '87. (1987)
- Super Cockpit Industry Days Briefing sponsored by the Air Force Systems Command, Human Systems Division, AAMRL. (March 31-April 1, 1987)
- Wellens, A. R., McNeese, M. D., "A Research Agenda for the Social Psychology of Intelligent Machines", Proceeding of NAECON '87. (1987)

ORIGINAL PAGE IS
OF POOR QUALITY